UNIVERSITY OF WASHINGTON
SEATTLE, WASHINGTON 98195

Department of Computer Science and Engineering   25 May 1996
Box 352350
206 685-1227 / FAX 206 543-2969
salesin@cs.washington.edu

Dr. Ralph Wachter
Program Manager/Officer ONR:311
Office of Naval Research
Ballston Tower One
800 North Quincy Street
Arlington, VA 22217-5660

Dear Dr. Wachter:

Here is a brief summary of some of the most significant advances during the past quarter, on ONR Young Investigator Award number N00014-95-1-0728:

Rendering of complex environments using a spatial hierarchy.[1] This paper, which will be presented at *Graphics Interface '96*, presents a new method for accelerating the rendering of complex static scenes. The technique is applicable to unstructured scenes containing arbitrary geometric primitives and has sublinear asymptotic complexity. Our approach is to construct a hierarchy of cells over the scene and to associate with each cell a simplified representation of its contents. The scene is then rendered using a traversal of the hierarchy in which a cell's approximation is drawn instead of its contents, if the approximation is sufficiently accurate. We apply the method to several different scenes and demonstrate significant speedups with little image degradation. We also exhibit and discuss some of the artifacts that our approximation may cause.

Rendering free-form surfaces in pen and ink.[2] In this paper, which will appear at SIG-GRAPH 96, we present new algorithms and techniques for rendering free-form surfaces in pen and ink. In particular, we introduce the concept of "controlled-density hatching" to convey tone, textures, and shapes. We also show how a planar map, a data structure central

---

[1]Rendering of complex environments using a spatial hierarchy, with B. Chamberlain, T. DeRose, D. Lischinski, and J. Snyder. To appear in *Graphics Interface '96*. An earlier version is available as Department of Computer Science and Engineering Technical Report TR 95-05-02, University of Washington, 1995.

[2]Rendering parametric surfaces in pen and ink, with G. Winkenbach. To appear in Proceedings of SIGGRAPH 96, in *Computer Graphics* Proceedings, Annual Conference Series, August 1996. Available as Department of Computer Science and Engineering Technical Report TR 96-01-05, University of Washington, 1996.

DEPARTMENT OF THE NAVY
OFFICE OF NAVAL RESEARCH
SEATTLE REGIONAL OFFICE
1107 NE 45TH STREET, SUITE 350
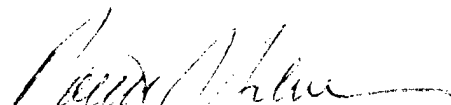SEATTLE WA 98105-4631

IN REPLY REFER TO:

4330
ONR 247
11 Jul 97

From: Director, Office of Naval Research, Seattle Regional Office, 1107 NE 45th St., Suite 350, Seattle, WA 98105
To: Defense Technical Center, Attn: P. Mawby, 8725 John J. Kingman Rd., Suite 0944, Ft. Belvoir, VA 22060-6218

Subj: RETURNED GRANTEE/CONTRACTOR TECHNICAL REPORTS

1. This confirms our conversations of 27 Feb 97 and 11 Jul 97. Enclosed are a number of technical reports which were returned to our agency for lack of clear distribution availability statement. This confirms that all reports are unclassified and are "APPROVED FOR PUBLIC RELEASE" with no restrictions.

2. Please contact me if you require additional information. My e-mail is *silverr@onr.navy.mil* and my phone is (206) 625-3196.

ROBERT J. SILVERMAN

to our rendering algorithm, can be constructed for curved surfaces. The planar map is used to clip strokes and to generate outlines. In addition, we show how to handle the casting of curved shadows onto curved objects. Finally, we introduce the use of traditional texture mapping techniques for controlling the tone of the illustration.

Reproducing color images as duotones.[3] In this paper, which will appear at SIGGRAPH 96, we investigate a new approach for reproducing color images. Rather than mapping the colors in an image onto the gamut of colors that can be printed with cyan, magenta, yellow, and black inks, we choose the set of printing inks for the particular image being reproduced. In this paper, we look at the special case of selecting inks for duotone printing, a relatively inexpensive process in which just two inks are used. Specifically, the system we describe takes an image as input, and allows a user to pre-select 0, 1, or 2 inks. It then chooses the remaining ink or inks so as to reproduce the image as accurately as possible and produces the apropriate color separations automatically.

Scale-dependent reproduction of pen-and-ink illustrations.[4] In this paper, which will appear at SIGGRAPH 96, we describe a compact resolution- and scale-independent representation for pen-and-ink illustrations. The proposed representation consists of a low-resolution grey-scale image, augmented by a set of discontinuity segments. We also present a new reconstruction algorithm that magnifies the low-resolution image while keeping the image sharp along the discontinuities. By storing pen-and-ink illustrations in this representation, we can produce high-fidelity illustrations at any scale and resolution by generating an image of the desired size and filling that image with pen-and-ink strokes.

Hierarchical image caching for accelerated walkthroughs of complex environments.[5] In this paper, which will appear at SIGGRAPH 96, we present a new method for accelerating walkthroughs of geometrically complex static scenes. As a preprocessing step, our method constructs a BSP-tree that hierarchically partitions the geometric primitives in the scene. In the course of a walkthrough, images of nodes at various levels of the hierarchy are cached for reuse in subsequent frames. A cached image is applied as a texture map to a single

---

[3]Reproducing color images as duotones, with J. L. Power, B. S. West, and E. J. Stollnitz. To appear in Proceedings of SIGGRAPH 96, in *Computer Graphics* Proceedings, Annual Conference Series, August 1996. Available as Department of Computer Science and Engineering Technical Report TR 96-01-08, University of Washington, 1996.

[4]Scale-dependent reproduction of pen-and-ink illustrations, with M. Salisbury, C. Anderson, and D. Lischinski. To appear in Proceedings of SIGGRAPH 96, in *Computer Graphics* Proceedings, Annual Conference Series, August 1996. Available as Department of Computer Science and Engineering Technical Report TR 96-01-02, University of Washington, 1996.

[5]Hierarchical image caching for accelerated walkthroughs of complex environments, with J. Shade, D. Lischinski, T. DeRose, and J. Snyder. To appear in Proceedings of SIGGRAPH 96, in *Computer Graphics* Proceedings, Annual Conference Series, August 1996. Available as Department of Computer Science and Engineering Technical Report TR 96-01-06, University of Washington, 1996.

quadrilateral that is drawn instead of the geometry contained in the corresponding node. Visual artifacts are kept under control by using an error metric that quantifies the discrepancy between the appearance of the geometry contained in a node and the cached image. The new method is shown to achieve significant speedups for a walkthrough of a complex outdoor scene, with little or no loss in rendering quality.

Interactive multiresolution surface viewing.[6] Multiresolution analysis has been proposed as a basic tool supporting compression, progressive transmission, and level-of-detail control of complex meshes in a unified and theoretically sound way. In this paper, which will be presented at SIGGRAPH 96, we extend previous work on multiresolution analysis of meshes in two ways. First, we show how to perform multiresolution analysis of colored meshes by separately analyzing shape and color. Second, we describe efficient algorithms and data structures that allow us to incrementally construct lower resolution approximations to colored meshes from the geometry and color wavelet coefficients at interactive rates. We have integrated these algorithms in a prototype mesh viewer that supports progressive transmission, dynamic display at a constant frame rate independent of machine characteristics and load, and interactive choice of tradeoff between levels of detail in geometry and color. The viewer operates as a helper application to Netscape, and can therefore be used to rapidly browse and display large collections of complex geometric models stored on the World Wide Web.

Multiresolution video.[7] In this paper, which will appear at SIGGRAPH 96, we present a new representation for time-varying image data, called multiresolution video. The representation allows for varying – and arbitrarily high – spatial and temporal resolutions in different parts of a video sequence. The representation is based on a sparse, hierarchical encoding of the video data. We show how multiresolution video supports a number of primitive operations: drawing frames at a particular spatial and temporal resolution; and translating, scaling, and compositing multiresolution sequences. These primitives are then used as the building blocks to support a variety of applications: video compression; multiresolution playback, including motion-blurred "fast-forward" and "reverse"; constant speed display; enhanced video scrubbing; and "video clip art" editing and compositing. The multiresolution representation requires little storage overhead, and the algorithms using the representation are both simple and efficient.

---

[6]Interactive multiresolution surface viewing, with A. Certain, J. Popović, T. DeRose, T. Duchamp, and W. Stuetzle. To appear in Proceedings of SIGGRAPH 96, in *Computer Graphics* Proceedings, Annual Conference Series, August 1996. Available as Department of Computer Science and Engineering Technical Report TR 96-01-07, University of Washington, 1996.

[7]Multiresolution video, with A. Finkelstein and C. Jacobs. To appear in Proceedings of SIGGRAPH 96, in *Computer Graphics* Proceedings, Annual Conference Series, August 1996. Available as Department of Computer Science and Engineering Technical Report TR 96-01-01, University of Washington, 1996.

The virtual cinematographer: a paradigm for automatic real-time camera control and directing.[8] In this paper, which will be presented at SIGGRAPH 96, we introduce a paradigm for automatically generating, in real-time, complete camera specifications for capturing events in virtual 3D environments. We demonstrate a fully implemented system called the Virtual Cinematographer as it is applied in a virtual "party" setting. The Virtual Cinematographer is implemented as a hierarchical finite state machine. Cinematographic expertise in the form of film idioms is encoded through a set of small finite state machines organized as a directed graph through call/return conventions and exception handling mechanisms. Each idiom is responsible for capturing a particular type of scene, such as three virtual actors conversing, or one actor moving across the environment. The idiom selects shot types and the timing of transitions between shots to best communicate events as they unfold. A set of camera modules shared by the idioms are responsible for the low level geometric placement of specific cameras for each shot type. The camera modules are also responsible for making subtle changes in the virtual actors' positions to best frame each shot.

Comic chat.[9] Comics have a rich visual vocabulary, and people find them appealing. They are also an effective form of communication. In this paper, which will be presented at SIGGRAPH 96, we describe a system, called Comic Chat, that represents on-line communications in the form of comics. Comic Chat automates numerous aspects of comics generation, including balloon construction and layout, the placement and orientation of comic characters, the default selection of character gestures and expressions, the incorporation of semantic panel elements, and the choice of zoom factor for the virtual camera. This paper describes the mechanisms that Comic Chat uses to perform this automation, as well as novel aspects of the program's user interface. Comic Chat is a working program, allowing groups of people to communicate over an on-line service. It has several advantages over other graphical chat programs, including the availability of a graphical history, and a dynamic graphical presentation.

Declarative camera control for automatic cinematography.[10] Animations generated by interactive 3D computer graphics applications are typically portrayed either from a particular character's point of view or from a small set of strategically-placed viewpoints. By ignoring camera placement, such applications fail to realize important storytelling capabilities that have been explored by cinematographers for many years. In this paper, which will ap-

---

[8] The virtual cinematographer: a paradigm for automatic real-time camera control and directing, with L. He and M. Cohen. To appear in Proceedings of SIGGRAPH 96, in *Computer Graphics* Proceedings, Annual Conference Series, August 1996.

[9] Comic chat, with D. Kurlander and T. Skelly. To appear in Proceedings of SIGGRAPH 96, in *Computer Graphics* Proceedings, Annual Conference Series, August 1996.

[10] Declarative camera control for automatic cinematography, with D. Christianson, S. E. Anderson, L. He, D. S. Weld, and M. F. Cohen. To appear in *AAAI '96*. An earlier version appeared as Department of Computer Science and Engineering Technical Report TR 95-01-03, University of Washington, 1995.

pear at AAAI '96, we describe several of the principles of cinematography and show how they can be formalized into a declarative language, called the "Declarative Camera Control Language" (DCCL). We describe the application of DCCL within the context of a simple interactive video game. Our prototype video game creates simulated animations in response to user input. Potential camera placements are generated by first segmenting the animations into individual scenes and then using a database of DCCL idioms to compile each scene into a set of possible films. Each potential film is then evaluated using domain-independent heuristics in order to select the best camera placements for filming the animation.

Global illumination of glossy environments using wavelets and importance.[11] In this paper, which recently appeared in *ACM Transactions on Graphics*, we show how importance-driven refinement and a wavelet basis can be combined to provide an efficient solution to the global illumination problem with glossy and diffuse reflections. Importance is used to focus the computation on the interactions having the greatest impact on the visible solution. Wavelets are used to provide an efficient representation of radiance, importance, and the transport operator. We discuss a number of choices that must be made when constructing a finite element algorithm for glossy global illumination. Our algorithm is based on the standard wavelet decomposition of the transport operator and makes use of a four-dimensional wavelet representation for spatially- and angularly-varying radiance distributions. We use a final gathering step to improve the visual quality of the solution. Features of our implementation include support for curved surfaces as well as texture-mapped anisotropic emission and reflection functions.

Clustering for glossy global illumination.[12] In this paper, submitted to *ACM Transactions on Graphics*, we present a new clustering algorithm for global illumination in complex environments. The new algorithm extends previous work on clustering for radiosity to allow for non-diffuse (glossy) reflectors. We treat each cluster as a point source with an angular distribution of radiant intensity, and we derive an error bound for transfers between these clusters. The algorithm groups input surfaces into a hierarchy of clusters, and then permits clusters to interact only if the error bound is below an acceptable tolerance. We show that the algorithm is asymptotically more efficient than previous clustering algorithms even when restricted to ideally diffuse environments. Finally, we demonstrate a working implementation of the algorithm on a complex architectural environment.

Wavelets for computer graphics. Wavelets are an exciting new mathematical tool, with roots in approximation theory, signal processing, and physics, that can be used for hierarchically
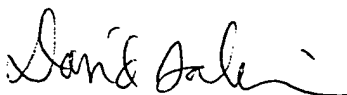
---

[11]Global illumination of glossy environments using wavelets and importance, with P. Christensen, E. Stollnitz, and T. DeRose. *ACM Transactions on Graphics* 15(1), January 1996.

[12]Clustering for glossy global illumination, with P. H. Christensen, D. Lischinski, and E. J. Stollnitz. Submitted to *ACM Transactions on Graphics*. An earlier version is available as Department of Computer Science and Engineering Technical Report TR 95-01-07, University of Washington, 1995.

decomposing functions. Wavelets have found a wide variety of applications in recent years, including signal analysis, image processing, and numerical analysis. Even more recently, wavelets are finding numerous applications in computer graphics, including image, curve, and surface editing and fast methods for solving simulation problems in global, as described in several of the projects above. To make the rather complex wavelets theory more accessible to the computer graphics community, Professor Tony DeRose, Eric Stollnitz (a graduate student), and I have written a primer that was published in two parts in *IEEE Computer Graphics and Applications*.[13] We recently extended this primer into a 250-page monograph, to be published this summer by Morgan-Kaufmann.[14]

I would be more than happy to furnish any of these papers, or discuss any of this work in more detail, upon your request.

Sincerely,

David Salesin
Assistant Professor

---

[13] Wavelets for computer graphics: A primer, with E. Stollnitz and T. DeRose. Published in two parts in *IEEE Computer Graphics and Applications* 15(3): 76–84 and 15(4): 75–85, May and July 1995. An earlier version is available as Department of Computer Science and Engineering Technical Report TR 94-09-11, University of Washington, 1994. A still earlier version appeared in the SIGGRAPH 94 "Computational Representations of Geometry" course notes #23, B. Naylor, editor, 113–141, 1994.

[14] *Wavelets for Computer Graphics: Theory and Applications*, with E. Stollnitz and T. DeRose. Morgan-Kaufmann Publishers, Inc., San Francisco, ISBN 1-55860-375-1. To appear in July, 1996.